

IN THE CLAIMS

Please amend the claims as follows:

1. (Currently Amended) A method of accessing shared memory in a computer system having a plurality of nodes, including a first node, wherein each node includes a processor and local memory, the method comprising:

distributing an application across the plurality of nodes;

building an application virtual address space, wherein building an application virtual address space includes:

building a local virtual address space for the application in each of the plurality of nodes, wherein:

the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node; and

the virtual address generated by the application executing on the node includes a node number of the node; and

exporting the local virtual address space for each node to a Remote Translation Table (RTT) associated with that node, wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation; and

performing a memory reference to a memory location in the application virtual address space, if the node number is not the local node number; wherein performing a memory reference to a memory location in the application virtual address space includes translating bits the node number of the application virtual address into a node address associated with the first node and translating bits of the application virtual address using the RTT associated with the first node.

2. (Original) The method of claim 1, wherein the local address space is read from a Translation Look-aside Buffer (TLB).

3. (Original) The method of claim 1, wherein building an application virtual address space further includes performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes to complete exporting their respective local virtual address spaces.
4. (Currently Amended) A system comprising:
a plurality of nodes, each node including:
one or more processors;
a memory; and
a memory controller operatively coupled to the memory and the one or more processors, wherein the memory controller includes a Remote Translation Table (RTT), wherein the RTT translates a virtual address received as part of a memory request received from another node into a memory request with physical addresses into the memory on the node associated with the RTT;
further wherein the RTT is initialized upon the start of a process associated with an application by building virtual to physical address translations for local virtual address space in the node corresponding to the application, wherein a virtual address includes a node number of the node, and exporting the virtual to physical address translations for the local virtual address space from the node to the Remote Translation Table (RTT) associated with that node, wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation.
5. (Original) The system of claim 4, wherein each of the plurality of nodes executes a synchronization operation that causes at least some of the plurality of nodes to wait for all of the plurality of nodes to complete exporting the virtual to physical address translations to their respective Remote Translation Tables.

6. (Currently Amended) A device-readable medium having instructions thereon that, when executed on a properly programmed information-processing device having a plurality of nodes, including a first node, each node having one or more processors, a memory, and a memory controller and coupled to the memory and the one or more processors, causes the information-processing device to perform a method comprising:

distributing an application across the plurality of nodes;

building an application virtual address space, wherein building an application virtual address space includes:

building a local virtual address space for the application in each of the plurality of nodes, wherein;

the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node; and

the virtual address generated by the application executing on the node includes a node number of the node; and

exporting the local virtual address space for each node to a Remote Translation Table (RTT) associated with that node, wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation; and

performing a memory reference to a memory location in the application virtual address space, if the node number is not the local node number, wherein performing a memory reference to a memory location in the application virtual address space includes translating bits the node number of the application virtual address into a node address associated with the first node and translating bits of the application virtual address using the RTT associated with the first node.

7. (Original) The device-readable medium of claim 6, wherein building a local virtual address space further includes performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes complete exporting their respective address space.

8. (Original) The device-readable medium of claim 6, wherein the local address space is read from a Translation Look-aside Buffer (TLB).

9. (Currently Amended) A multinode system for implementing remote address translation, the system comprising:

a plurality of nodes, including a first node, each of the plurality of nodes including:

one or more processors,

a memory, and

a memory controller operatively coupled to the memory and the one or more processors;

means for distributing an application across the plurality of nodes;

means for building an application virtual address space, wherein the means for building an application virtual address space includes:

means for building a local virtual address space for the application in each of the plurality of nodes, wherein;

the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node; and

the virtual address generated by the application executing on the node includes a node number of the node; and

means for exporting the local virtual address space for each node to a Remote Translation Table (RTT) associated with that node, wherein the means for exporting includes means, within each node, for requesting that the operating system load the RTT from the local address space of its respective node and means for requesting that the operating system enable remote translation; and

means for performing a memory reference to a memory location in the application virtual address space, if the node number is not the local node number, wherein performing a memory reference to a memory location in the application virtual address space includes:

means for translating bits the node number of the application virtual address into a node address associated with the first node, and

means for translating bits of the application virtual address using the RTT associated with the first node.

10. (Original) The multinode system of claim 9, wherein building an application virtual address space further includes a means for performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes to complete exporting their respective local virtual address spaces.

11. (Currently Amended) A multi-node system for implementing remote address translation, the system comprising:

a network;

a source node coupled to the network, wherein the source node includes a first remote-translation table (RTT);

a remote node coupled to the network, wherein the remote node includes a second RTT; wherein on the remote node the second RTT is built using a first local address space on the source node exported from the source node to the remote node using an operating system call to perform the export;

wherein on the source node the first RTT is built using a second local address space on the remote node exported from the remote node to the source node using the operating system call to perform the export; and

the operating system enables remote translation utilizing the first and second RTTs; and wherein both the first and second RTTs include one or more virtual address and each virtual address includes a node number of a remote node that built the virtual address.

12. (Previously Presented) The method of claim 1, wherein requesting the operating system enable remote translation passes control of the RTT to the operating system.

13. (Previously Presented) The method of claim 12, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

14. (Previously Presented) The system of claim 4, wherein requesting the operating system enable remote translation passes control of the RTT to the operating system.
15. (Previously Presented) The system of claim 14, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.
16. (Previously Presented) The device-readable medium of claim 6, wherein requesting the operating system enable remote translation passes control of the RTT to the operating system.
17. (Previously Presented) The device-readable medium of claim 16, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.